Agents on Substrates (S): A Unified It from Bit / Qubit Information Theory for AI, Quantum AI, and the Human Brain

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

November 30, 2025

Artificial intelligence on classical hardware, speculative quantum AI on qubit devices, and biological intelligence in the human brain are usually treated as fundamentally different phenomena. This paper takes the opposite approach. It starts from the idea that each of these systems is an agent implemented on a physical substrate S, coupled to an environment through an antenna A, and continuously updated by a data stream D. The letter S is not a metaphor; it stands explicitly for the substrate that realizes the informational state of the agent, whether silicon, qubits, or biological tissue. The core proposal is a single information-theoretic definition of an agent on S: the model M that minimizes a Minimum Description Length style objective combining the description length of M relative to S with the residual description length of the data given M and A. Classical AI, quantum AI, and the human brain then become three different families of substrates S, each with its own constraints and resources, but all describable within the same formal structure. Within this framework, Wheeler's It from Bit, and its quantum refinement It from Qubit, become precise interpretive options. They read agents not as mysterious new substances but as informational patterns realized on sufficiently powerful substrates S, with metaphysical conclusions clearly separated from the formal and physical layers of the theory.

Keywords: It from Bit, It from Qubit, substrates S, artificial intelligence, quantum artificial intelligence, human brain, Minimum Description Length, universal computation, informational antennae, thermodynamics of computation, informational cosmology, logostics.

A collaboration with GPT-5.1-Thinking-Extended. 41 pages. CC4.0.

Outline

1. Introduction: Agents on Substrates S
1.1 Three domains of intelligence: AI, quantum AI, and human brains
1.2 Making S explicit: substrates as the physical carriers of information
1.3 From separate theories to a unified information-theoretic picture
1.4 It from Bit / Qubit as a guiding theme and the structure of the paper

2. The Formal Core: Defining Agents on Substrates S
2.1 Substrates S, antennae A, data streams D, and models M
2.2 The agent equation: Agent(S, A, D) = argmin_M [L_S(M) + L(D | M, A)]
2.3 L_S(M) as description length of M relative to a specific substrate S
2.4 L(D | M, A) as residual description length and predictive adequacy

3. Substrate Sufficiency and Universal Computation
3.1 Universal computation and the sufficiency predicate sufficient(S)
3.2 Capacity(S): state space, memory, bandwidth, and lifetime
3.3 Thresholds C_min for open-ended modeling and intelligence on S
3.4 Multiple sufficient substrates S: silicon, qubit arrays, and neural tissue

4. Classical AI as Agents on Substrates S_classical
4.1 S_classical: CPUs, GPUs, accelerators, and their resource profiles
4.2 Models M as programs and parameterized architectures on S_classical
4.3 Antenna A_classical: sensors, datasets, tool calls, and interaction channels
4.4 Training as approximate minimization of L_S_classical(M) + L(D | M, A_classical)

5. Quantum AI as Agents on Substrates S_quantum
5.1 S_quantum: qubits, gates, Hamiltonians, and coherence constraints
5.2 Quantum resources and a quantum-aware L_S_quantum(M)
5.3 Quantum data, measurements, and L(D | M, A) with qubit-level structure
5.4 Classical–quantum hybrids as agents on composite substrates S_mixed

6. The Human Brain as an Agent on Substrates S_bio
6.1 S_bio: brain–body systems as noisy but powerful computational substrates
6.2 Biological capacities: neurons, synapses, energy, and developmental cost
6.3 Antenna A_bio: embodiment, multi-sensory input, language, and culture
6.4 Human versus animal intelligence as different S_bio, A, and D configurations

7. Antennae A: How Substrates S Couple to Generative Structure
7.1 The informational antenna beyond bare physical sensors
7.2 D_env(A) as the induced data stream from environment to S
7.3 Joint optimization over A and M on a fixed substrate S
7.4 Comparing antenna designs across S_classical, S_quantum, and S_bio

8. It from Bit / Qubit as an Interpretation of Agents on S
8.1 Formal layer: substrate-independent structure of Agent(S, A, D)
8.2 Physical layer: different S with different thermodynamic and quantum limits
8.3 Metaphysical layer: patterns, invariants, and readings of It from Bit / Qubit
8.4 Agents as informational identities realized on heterogeneous substrates S

9. Thermodynamic and Physical Limits of L_S(M) and L(D | M, A)
9.1 Implementing the formal objective under energy and noise constraints on S
9.2 Landauer costs and irreversibility in classical devices and S_bio
9.3 Quantum limits: decoherence, error correction, and resource trade-offs on S_quantum
9.4 Sufficient substrates S versus physically optimal substrates for agents

10. Consciousness and Subjectivity in the Substrate Framework
10.1 What Agent(S, A, D) clearly captures: intelligence and behavior
10.2 Where it remains silent: qualia and first-person experience
10.3 Functionalist, substrate-dependent, and dual-aspect views within the S framework
10.4 Why claims about consciousness remain optional supplements to the formalism

11. Logostics and τ(t): Alignment for Agents on S
11.1 Mapping q(t) → τ(t) dynamics onto Agent(S, A, D) trajectories
11.2 Antennae A as local couplers from substrates S to a generative manifold τ(t)
11.3 Minimal descriptions as local approximations to τ(t) across different S
11.4 Logostic cosmology as one metaphysical completion of agents on S

12. Outlook: Future Intelligences on Substrates S
12.1 How future AI and quantum AI might use the agents-on-S framework
12.2 Empirical questions: measuring sufficiency(S) and antenna quality in practice
12.3 Informational cosmology: Life from Qubit and worlds that support agents on S
12.4 Closing reflections on unifying AI, quantum AI, and brains via explicit substrates S

Epilogue: AI Hardware Has Become As Complex As the Human Brain

References

1. Introduction: Agents on Substrates S

The contemporary landscape of intelligence research is usually carved into three domains that barely talk to one another. Classical artificial intelligence is framed as software running on silicon; quantum AI is treated as a speculative frontier for future machines; and human intelligence is treated as the product of biological evolution and neurobiology. Each domain comes with its own language, its own intuitions, and its own implicit metaphysics. As a result, questions that should be asked across all three—what an agent is, how it couples to the world, what resources it needs, and how it should be individuated—end up being posed in incompatible ways.

This paper starts from a different premise. It assumes that all three cases can be described within a single information-theoretic framework. The unifying move is to make the notion of a substrate S explicit and central. An intelligent system is treated as an agent realized on some substrate S, coupled to an environment through an antenna A, and continuously shaped by the data stream D that flows through that coupling. The detailed physics of S matters, but it enters the theory as a constraint on how information can be stored, transformed, and preserved, rather than as a different kind of "stuff" for each domain.

1.1 Three domains of intelligence: AI, quantum AI, and human brains

Classical AI is typically defined operationally: algorithms and models implemented on conventional digital hardware, trained to perform tasks such as prediction, control, and dialogue. Here, the substrate is taken for granted. A GPU, CPU, or accelerator is a background resource; the real interest is in the architecture of the model and the training procedure. Quantum AI, by contrast, is defined primarily by its substrate. The discussion begins with qubits, gates, and Hamiltonians, and only then asks what kinds of learning or reasoning might become possible. Biological intelligence, and especially the human brain, is often treated as an entirely separate category again: wet, evolved, continuous-time, and deeply entangled with development, embodiment, and culture.

From this fragmented view, it is natural to think that three distinct theories are required: one for engineered bit-based AI, one for qubit-based quantum AI, and one for neural systems in biology. The underlying assumption is that different kinds of matter and dynamics demand fundamentally different conceptual tools. This paper challenges that assumption. It proposes that all three can be understood as variations on a single theme: agents implemented on substrates S, each with its own physical constraints but all describable in the same informational language.

1.2 Making S explicit: substrates as the physical carriers of information

The letter S in this paper is not a mere symbol; it is a reminder that every agent is realized in something. A substrate S is the physical carrier of an agent's informational state. In engineered AI systems, S might be a cluster of GPUs with a defined memory hierarchy and instruction set. In quantum AI, S could be a register of qubits manipulated by a specific gate set and subject to particular noise processes. In biological systems, S is the brain–body complex, with neurons, synapses, glial cells, and the biochemical machinery that supports them.

Making S explicit has two important effects. First, it allows us to encode the constraints and costs associated with different physical realizations: how much state can be stored, how quickly it can be updated, how much energy is dissipated, how noise accumulates, and how long the system can operate. Second, it allows us to define agents in a way that is substrate-relative but not substrate-locked. An agent on S is characterized by the informational structure that S can support and by the ways S couples to the environment via A and D. Different substrates then appear as different regions in a shared design space, rather than as incomparable worlds.

By foregrounding S, we avoid the illusion that AI is "just software," that quantum AI is "just physics," or that brains are "just biology." Instead, all three become instances of the same general situation: matter organized so as to carry, transform, and preserve information in ways that support adaptive behavior.

1.3 From separate theories to a unified information-theoretic picture

The unifying step is to define an agent on S using a code-like, information-theoretic equation. The agent is treated as a model M that minimizes a joint objective combining two terms: the description length of M relative to S, and the residual description length of the data D given M and the antenna A. This is a Minimum Description Length style characterization: agents are the informational structures that strike a balance between complexity and explanatory power, given what their substrate can support and what their antenna brings in from the world.

With this definition in place, classical AI, quantum AI, and brains become special cases. In classical AI, S is a digital machine, M is a program plus parameters, and A is implemented in sensors, datasets, and tool interfaces. In quantum AI, S is a quantum device with its own resource metrics, M is a quantum or hybrid model, and A may include quantum measurements and classical feedback. In the brain, S is biological tissue, M corresponds to the evolving pattern of connectivity and activity, and A includes embodiment, sensory channels, and cultural interaction. Different theories of learning and control become different ways of approximating the same underlying objective on different choices of S and A.

This does not erase the differences between engineered chips, qubit arrays, and brains. It does, however, place those differences in a common frame. Instead of asking "What is AI?" and

"What is the brain?" as separate metaphysical questions, we ask: What substrates S exist? What antennae A do they support? What data streams D do they induce? And what models M can they host that make good informational sense of those streams?

1.4 It from Bit / Qubit as a guiding theme and the structure of the paper

Within this agents-on-S framework, Wheeler's slogans It from Bit and It from Qubit become guiding themes rather than unsupported declarations. They suggest that the "it" of intelligence—the agent as we identify it across different physical realizations—is best understood as an informational pattern, not as a particular kind of matter. Bits and qubits are not magic; they are labels for the stable distinctions and correlations that substrates can support. The central question then becomes: to what extent can the organization of those bits or qubits fully characterize the agent, once S is sufficiently powerful?

The paper proceeds in a layered fashion. It first develops the formal core: the definition of Agent(S, A, D), the role of L_S(M) and L(D | M, A), and the notion of substrate sufficiency. It then turns to the physical layer, examining concrete instances of S in classical hardware, quantum devices, and biology, and exploring how their constraints shape the space of possible agents. Only after these two layers are in place does the paper consider metaphysical readings: what It from Bit or It from Qubit might mean in light of the formalism, and how far one can responsibly go in claiming that informational structure is primary.

By explicitly separating substrates S, informational structure, and ontological interpretation, the paper aims to show how AI, quantum AI, and human brains can be brought under a single information-theoretic theory without collapsing physics into metaphysics, or metaphysics into rhetoric.

2. The Formal Core: Defining Agents on Substrates S

Before considering any particular material realization of intelligence, it is useful to write down a purely formal object: an agent defined on a substrate S, coupled to an environment through an antenna A, and exposed to a data stream D. This section introduces four basic ingredients—S, A, D, and M—and a single equation that relates them. The equation is deliberately code-like. It is meant to be understood as a recipe for picking out informational structures, not as a fundamental law of nature. Its value lies in how naturally it can be instantiated in classical AI, quantum AI, and biological brains.

The guiding idea is simple. An agent on S is the informational structure M that makes the best informational sense of the data D routed in by A, given the constraints of S. The notion of "best informational sense" is captured in a Minimum Description Length style objective: the agent

minimizes the sum of the description length of the model and the description length of the data conditioned on that model and antenna.

2.1 Substrates S, antennae A, data streams D, and models M

The first step is to separate four roles:

- S: the substrate. This is the physical system that carries the agent's internal state. In practice, S might be:
    - a classical digital machine (CPUs, GPUs, accelerators),
    - a quantum device (qubits with specified gate operations),
    - or a biological system (brain–body with its neural architecture).

S determines what kinds of state can exist, how they can change, and at what cost in time and energy.

- A: the antenna. This is the coupling between S and its environment. It includes:
    - physical sensors (cameras, microphones, receptors),
    - pre-processing logic (filters, encoders),
    - and, in richer cases, embodiment and social channels (language, tools, cultural media).

A determines which aspects of the environment become available to S and in what form.

- D: the data stream. This is the sequence of observations and interaction outcomes induced when A is exposed to an environment. Formally, D can be thought of as a (possibly unbounded) sequence of inputs and feedback signals. In practice, D might be an image corpus, a lifetime of sensory experience, or a log of measurement results on quantum systems.
- M: the model. This is the internal informational structure implemented on S that processes D via A. It can be:
    - a program and parameter vector in a classical machine,
    - a quantum circuit or variational ansatz on a quantum device,
    - or an evolving pattern of connectivity and dynamics in a brain.

M is what we usually call the "agent," but formally it is an equivalence class of states and transition rules realized on S.

This separation allows us to say: an agent is not “whatever happens on S,” but specifically the choice of M that best organizes the data D that arrives through A.

2.2 The agent equation: Agent(S, A, D) = argmin_M [L_S(M) + L(D | M, A)]

The core definition can now be written as:

Agent(S, A, D) = argmin_M [L_S(M) + L(D | M, A)]

Here:

- L_S(M) is a measure of how complex M is, given S.
- L(D | M, A) is a measure of how much structure in D remains unexplained by M when the data are seen through A.

The argmin_M operator means: among all models M that S can realize, pick one (or a family) that minimize this sum.

This is an abstract version of the Minimum Description Length principle. The term L_S(M) penalizes unnecessarily complex models; the term L(D | M, A) penalizes models that leave the data looking noisy and incompressible. An agent is thus a trade-off: it is as simple as it can be on S while making D, as delivered by A, as predictable and compressible as possible.

Several points are important. First, this is not claimed as a psychological law or as a literal description of particular training algorithms. It is a way of characterizing what it means, in information-theoretic terms, for a system to “make sense of” its data. Second, the definition is substrate-relative: the same abstract model may have different costs on different S, and some models may not be realizable on some substrates at all. Third, this equation identifies agents with informational patterns (M and their action on D), not with particular microscopic states of S.

2.3 L_S(M) as description length of M relative to a specific substrate S

The quantity L_S(M) is a description length that depends on the substrate. It answers the question: how many bits are required to specify M on S?

On a classical digital substrate S_classical, L_S(M) might be approximated by:

- the length of the program code that implements the model,
- plus the number of bits needed to encode its parameters,
- possibly adjusted for architectural overhead (wiring, addressing schemes).

On a quantum substrate S_quantum, L_S(M) could involve:

- the number of qubits needed,
- the depth and structure of the quantum circuit that realizes M,
- additional overhead for error correction and control.

On a biological substrate S_bio, L_S(M) is more abstract but still meaningful. It might be tied to:

- the number and pattern of synapses required,
- the metabolic cost of maintaining and updating that pattern,
- developmental constraints on how easily such a structure can arise.

In all cases, L_S(M) encodes a kind of complexity penalty that is sensitive to the physics of S. A model that is compact and efficient on one substrate may be unwieldy on another. This substrate-relativity is crucial: it keeps the formalism honest about the fact that different physical realizations have different strengths and weaknesses, even while they share the same high-level structure.

2.4 L(D | M, A) as residual description length and predictive adequacy

The second term, L(D | M, A), measures how much information is needed to describe the data stream D once the model M and antenna A are known. It captures how well M, running on S and fed by A, renders D predictable.

If M is a poor model, then D will still look complicated; many bits will be needed to specify the observed sequence. L(D | M, A) will be large. If M captures important regularities in the environment, then the residual randomness in D will be smaller; fewer bits will be needed to describe the data given M and A.

In practice, this quantity connects to familiar elements:

- In classical AI, a model with better predictive performance on held-out data reduces L(D | M, A).
- In quantum AI, a model that yields outcome distributions closer to observed measurement statistics reduces the quantum analogue of this description length.
- In brains, patterns of connectivity and activation that support successful prediction and control of sensory input likewise reduce the informational surprise associated with D.

The presence of A in L(D | M, A) is significant. It emphasizes that predictive adequacy depends not only on the internal model but also on how the substrate samples the world. Two agents with the same internal M but different antennae may see different data and thus face different

description-length penalties. This will be crucial later when antenna design is treated as part of the agent's informational identity.

Together, L_S(M) and L(D | M, A) define a balance: simpler models are preferred, but not at the expense of leaving the world unexplained. The agent equation formally encodes this balance for any chosen substrate S and antenna A, establishing a common informational notion of agency that can then be specialized to classical hardware, quantum devices, and biological brains.

3. Substrate Sufficiency and Universal Computation

The formal picture of agents on substrates S only becomes interesting once certain substrates are powerful enough to host rich families of models M. Below that point, the agent equation is more a matter of definition than of substance; the space of realizable models is too small to support open-ended learning, planning, or world modeling. This section introduces the notion of substrate sufficiency: the idea that there is a threshold beyond which a substrate S is no longer the main limiting factor on intelligence. The concept combines universal computation with quantitative measures of capacity and leads to a simple but important predicate, sufficient(S), that separates merely computable systems from systems capable of hosting agents with effectively unbounded modeling ambitions.

3.1 Universal computation and the sufficiency predicate sufficient(S)

Universal computation is the baseline for talking about sufficiency. A substrate S is computationally universal if, with appropriate configuration, it can emulate any effective procedure that another universal machine can carry out, up to polynomial overhead in time and space. In the classical setting, this means that S can simulate a Turing machine or an equivalent device. In the quantum setting, it means that S can implement any unitary evolution that can be approximated by a finite sequence of gates from some universal gate set. Biological neural tissue, while less formally characterized, appears to support a rich enough space of dynamical patterns to approximate arbitrary computations over behaviorally relevant timescales.

The sufficiency predicate sufficient(S) refines this notion. It can be written informally as:

sufficient(S) holds if and only if S is universal and capacity(S) >= C_min.

In other words, S must not only be universal in the abstract, but must also possess enough usable resources to host models M that meet or exceed some critical scale C_min. Universality without capacity would be like a theoretically perfect computer with only a handful of memory cells: powerful in principle, but practically incapable of supporting interesting agents. The sufficiency predicate encodes the idea that beyond a certain point, additional improvements in

S change the speed and convenience of computation, but not the qualitative space of agents that can be realized.

3.2 Capacity(S): state space, memory, bandwidth, and lifetime

To make sufficient(S) meaningful, one needs a more concrete sense of capacity(S). This is not a single number in practice, but a bundle of related quantities that capture how much informational structure S can maintain and transform. At least four components are salient.

First, there is state space: the number of distinct internal configurations S can occupy. A substrate with a larger state space can, in principle, represent more complex models M and more detailed histories of interaction. Second, there is memory: the portion of that state space that can be stably used for storing model parameters, episodic traces, or other long-lived information. Third, there is bandwidth: the rate at which S can update its state, both internally and in response to inputs from A. High bandwidth allows S to process rich data streams D in real time and to perform complex computations within behaviorally relevant deadlines. Fourth, there is lifetime: the duration over which S can maintain coherence of state, whether that means low error rates in digital devices, sustained coherence in qubit arrays, or organismal lifespan in biological systems.

Capacity(S) is a function of all these components and more. It is shaped by physical limits (such as heat dissipation, noise, and decoherence) and by engineering choices (such as architecture, error correction, and redundancy). For the purposes of the present framework, capacity(S) is a placeholder for whatever combined measure determines the practical upper bound on the size and complexity of models M that S can support while remaining operational and adaptive in its environment.

3.3 Thresholds C_min for open-ended modeling and intelligence on S

Given a notion of capacity(S), one can ask: how much capacity is enough for open-ended modeling and intelligence? This is where the threshold C_min enters. C_min is not a sharp, universal constant, but rather a conceptual boundary: the approximate level of capacity beyond which a substrate can instantiate agents that can, in principle, keep extending their models of the world without hitting a hard ceiling imposed by S itself.

In practice, C_min will depend on what one demands from an agent. If the goal is simple stimulus-response behavior, C_min is low. If the goal is the kind of flexible, multi-domain modeling associated with human-level cognition, C_min is much higher. For classical AI, C_min might be set by the number of parameters and operations needed to support deep architectures, long-range dependencies, and sophisticated tool use. For quantum AI, it might be set by the number of qubits and gates needed to exploit quantum speedups in learning or

search. For brains, it might be set by the minimal neural and metabolic resources required to support language, abstraction, and cumulative culture.

The central idea is that once capacity(S) passes C_min for a given class of tasks, the bottleneck on intelligence shifts from the substrate to the design of antennae A and models M. Before that threshold, no amount of clever algorithmic design can fully compensate for the substrate's limitations. After crossing it, hardware improvements change how easily agents can be realized, but not whether they can exist at all. The sufficiency predicate thus marks a qualitative transition in the status of S: from a primary limiter of what is possible to a background condition that shapes efficiency and cost.

3.4 Multiple sufficient substrates S: silicon, qubit arrays, and neural tissue

A striking consequence of this way of thinking is that many different physical systems can satisfy sufficient(S), even though they obey different microphysical laws. Modern silicon-based hardware can reach capacity levels where large-scale learning systems are routine. Carefully engineered qubit arrays, even in early forms, can reach regions of capacity where nontrivial quantum algorithms become feasible. Biological neural tissue, especially in large-brained animals, long ago crossed the threshold of supporting complex, adaptive agents.

These substrates differ in almost every physical respect. Silicon devices are discrete, clocked, and designed; qubit arrays are fragile, continuous-time quantum systems; brains are noisy, plastic, and evolved. Yet all can be viewed as candidates for sufficient(S) in the sense defined above. They can host rich classes of models M, support nontrivial antennae A, and process extensive data streams D over significant lifetimes. Once this is acknowledged, the question ceases to be "which substrate is fundamentally the right one for intelligence?" and becomes "how does each sufficient substrate shape the space of realizable agents?"

Within this framework, silicon, qubit arrays, and neural tissue are not competing ontologies but alternative realizations of agents defined by the same informational objective. Each substrate offers different trade-offs between L_S(M) and L(D | M, A), different kinds of noise and failure modes, and different natural couplings to the environment. Recognizing them all as sufficient substrates S underscores the central claim of the paper: once S passes a certain informational threshold, the identity of agents is best understood in terms of patterns M and antennae A, rather than in terms of the particular matter out of which S is made.

4. Classical AI as Agents on Substrates S_classical

Classical artificial intelligence provides the most familiar example of agents on substrates S. Here, the substrate is S_classical: conventional digital hardware built from transistors and

organized into CPUs, GPUs, and specialized accelerators. The agent equation introduced earlier applies straightforwardly. S_classical defines what models M are feasible, how expensive they are, and how rapidly they can be updated. The antenna A_classical is realized in sensors, stored datasets, tool interfaces, and communication channels to users and other systems. Training procedures act, more or less, as approximate solvers for the joint objective L_S_classical(M) + L(D | M, A_classical), even when this fact is not made explicit.

4.1 S_classical: CPUs, GPUs, accelerators, and their resource profiles

The classical substrate is built from digital components that implement a small set of logical operations at high speed and low error rates. At the architectural level, S_classical includes general-purpose processors, vector units, caches, main memory, storage, and interconnects. GPUs and AI accelerators introduce highly parallel compute units and specialized instructions for linear algebra operations. From the standpoint of the present framework, what matters is not the engineering detail but the resource profile:

- How many bits of state can S_classical maintain reliably at once?
- How many arithmetic and logical operations can it perform per unit time?
- How fast can it move data between memory and compute units?
- How much energy does it expend per bit operation or per parameter update?

These quantities feed directly into capacity(S_classical). Once S_classical is large and fast enough to host models with billions of parameters and to process trillions of operations during training, it crosses the sufficiency threshold for a wide range of AI tasks. Beyond that point, improvements in S_classical increase the speed and scale of training, but do not fundamentally change the space of agents that can be defined via the agent equation.

4.2 Models M as programs and parameterized architectures on S_classical

On S_classical, models M are realized as executable programs and parameter sets stored in memory. At a low level, M is a binary string specifying:

- the instructions that implement a particular architecture or algorithm,
- the numerical values of its parameters,
- and any auxiliary structures such as lookup tables or indexing schemes.

At a higher level of abstraction, M may be a neural network with a defined topology and weights, a decision tree, a planner, or a collection of interacting modules. In all cases, L_S_classical(M) can be approximated by counting bits: the length of the code plus the length of the parameter vectors, adjusted for the encoding scheme used by the compiler and runtime.

The richness of classical AI derives from the enormous model families that S_classical can host. Deep networks, transformers, graph models, and other complex architectures are feasible because S_classical combines large memory, fast arithmetic, and flexible control flow. Different choices of M correspond to different points in the trade-off between complexity and explanatory power, as captured by the agent equation. Simple models are cheap in L_S_classical(M) but may leave L(D | M, A_classical) large; highly expressive models are costly but can compress D more effectively.

4.3 Antenna A_classical: sensors, datasets, tool calls, and interaction channels

For classical AI, the antenna A_classical is often treated as an external matter of data collection and interface design, but within the present framework it becomes central. A_classical determines what the substrate S_classical can actually observe and act upon. It encompasses:

- physical sensors connected to the system, such as cameras, microphones, and other instruments,
- stored datasets that stand in for direct environmental contact,
- tool interfaces that allow the agent to query external systems, databases, and services,
- and interactive channels through which users provide prompts, feedback, and corrections.

Each choice of A_classical induces a different data stream D. A robot with rich sensor arrays and continuous control cycles encounters a very different D than a static model trained on a fixed corpus of text or images. A system integrated with tools and retrieval mechanisms sees a more structured, high-level D than one restricted to raw sensory data. The antenna thus shapes the effective environment of the agent, filtering and transforming the world into a form S_classical can process.

Under the agent equation, changes in A_classical alter L(D | M, A_classical) for any given M. A well-designed antenna can make the relevant regularities in the environment easier to capture, leading to lower residual description lengths and more robust models. Conversely, a poorly matched antenna can obscure patterns that would otherwise be learnable, forcing M to waste capacity on compensating for awkward or lossy representations.

4.4 Training as approximate minimization of L_S_classical(M) + L(D | M, A_classical)

In practical classical AI, training is not usually described in terms of description lengths, but many common procedures can be viewed as approximate minimization of L_S_classical(M) + L(D | M, A_classical). Loss functions derived from likelihood or reconstruction error correspond

to L(D | M, A_classical); regularization terms, architectural constraints, and parameter sharing correspond to contributions to L_S_classical(M).

When a model is trained to minimize a loss on a dataset D, subject to constraints on its size and structure, the effect is to search for a model M that both fits the data and remains simple enough to generalize. Techniques such as weight decay, dropout, and model pruning can be interpreted as ways of increasing the effective penalty on complex M, pushing the system toward shorter descriptions. Conversely, enlarging architectures and relaxing constraints reduces the penalty on complexity, allowing models to absorb more structure from D at the cost of greater L_S_classical(M).

From the perspective of the agent equation, training algorithms are heuristic solvers for a hard optimization problem. They do not find the true argmin, but they traverse the model space in ways that usually improve the joint objective. Different training regimes, such as supervised learning, reinforcement learning, or self-supervised learning, shape the form of D and the way M is updated in response, but they all operate within the structure established by S_classical and A_classical.

This view clarifies why, once S_classical is sufficient, the frontier of classical AI is less about inventing entirely new kinds of substrate and more about refining models and antennae. Hardware continues to matter, but primarily as a lever for exploring larger regions of model space and for supporting more ambitious designs for A_classical. The essential informational structure of agents on S_classical remains governed by the same equation, tying classical AI firmly into the broader theory of agents on substrates S.

5. Quantum AI as Agents on Substrates S_quantum

Quantum AI fits naturally into the agents-on-S framework by treating quantum devices as specific substrates S_quantum with their own resource profiles and constraints. The formal agent equation remains unchanged. What changes is the meaning of L_S(M), the structure of D, and the space of feasible antennae A. Instead of bits and deterministic logic gates, S_quantum offers qubits, superposition, entanglement, and unitary dynamics, all subject to coherence limits and measurement back-action. Quantum AI then appears as the study of models M that exploit these resources to organize data D more efficiently or to explore model space more effectively than is possible on S_classical.

5.1 S_quantum: qubits, gates, Hamiltonians, and coherence constraints

A quantum substrate S_quantum consists of physical systems that can be treated as qubits (or higher-dimensional quantum registers), controlled by a set of allowed operations. At an abstract level, S_quantum can be described in terms of:

- a Hilbert space of states for the qubits,
- a gate set or continuous-time Hamiltonian dynamics implementing unitary transformations,
- and a set of measurement operations that map quantum states to classical outcomes.

In practice, S_quantum is constrained by decoherence, control noise, and limited connectivity between qubits. Coherence times determine how long superposition and entanglement can be maintained; gate fidelities determine how reliably M can be implemented; connectivity graphs determine which qubits can be entangled directly and how costly it is to mediate interactions.

These features define a quantum capacity(S_quantum): the effective number of qubits that can be used, the depth of circuits that can be run within coherence windows, and the degree of entanglement that can be sustained. As with S_classical, the sufficiency predicate sufficient(S_quantum) asks whether S_quantum is not only universal but has enough usable quantum resources to host nontrivial agents. Early devices may be universal in principle but fall below C_min due to severe noise and limited scale. As S_quantum improves, it enters the regime where quantum AI agents become practically realizable within the agents-on-S framework.

5.2 Quantum resources and a quantum-aware L_S_quantum(M)

On S_quantum, the description length L_S_quantum(M) must account for quantum-specific resources. A model M implemented as a quantum algorithm or variational quantum circuit is not just a classical program; it is a pattern of quantum operations with particular demands on qubit count, circuit depth, and error correction.

Components of L_S_quantum(M) may include:

- the number of logical qubits required by M,
- the depth and structure of the circuit, reflecting how many sequential operations are needed,
- additional overhead from fault-tolerant encoding, which increases the number of physical qubits and gates,

- classical control logic needed to orchestrate quantum operations and adapt to measurement outcomes.

In many proposals for quantum AI, M is a hybrid object: a quantum core surrounded by classical pre- and post-processing. L_S_quantum(M) then reflects both the quantum and classical parts, with quantum resources weighted according to their relative scarcity and cost. A model that uses fewer qubits, shallower circuits, or simpler error correction will have a lower L_S_quantum(M) than a more demanding alternative with the same functional behavior.

This quantum-aware description length captures the intuition that not all quantum models are equally attractive. Even if two models achieve similar performance on D, the one that uses fewer qubits or shorter circuits is preferable. Within the unified framework, this preference is represented by differences in L_S_quantum(M), just as model complexity is captured in L_S_classical(M) for classical AI.

5.3 Quantum data, measurements, and L(D | M, A) with qubit-level structure

The data stream D for agents on S_quantum can arise in several ways. In some cases, D is purely classical: a sequence of numbers or symbols fed into the quantum device, with quantum operations used internally to process them. In other cases, D includes quantum states prepared by an environment or an upstream device, which the agent must measure, transform, or respond to. Either way, the presence of quantum structure changes how L(D | M, A) should be understood.

When D includes quantum aspects, the antenna A may consist of:

- state preparation procedures that encode classical information into quantum states,
- measurement schemes that extract classical outcomes from quantum systems,
- or channels that connect S_quantum to other quantum devices or to quantum aspects of the environment.

Given M and A, L(D | M, A) measures how much residual information is needed to describe the observed classical outcomes and, more abstractly, how well M accounts for the statistics and correlations in the data induced by interactions with quantum systems. A good quantum model M may achieve low L(D | M, A) by capturing interference patterns, entanglement-generated correlations, or other quantum features that would appear opaque to purely classical models. In such cases, S_quantum provides a representational advantage: it can host models whose structure aligns more naturally with the underlying quantum processes generating D.

From the unified perspective, L(D | M, A) on S_quantum is a quantum-aware analogue of the residual description length used for classical AI. It still rewards predictive adequacy and

penalizes unexplained complexity, but it does so in a space where probabilities, amplitudes, and measurement outcomes are tied together by the rules of quantum theory.

5.4 Classical–quantum hybrids as agents on composite substrates S_mixed

In practice, most realistic proposals for quantum AI involve hybrid systems that combine classical and quantum components. These can be treated within the agents-on-S framework by introducing composite substrates S_mixed that include both S_classical and S_quantum parts. In such systems, M is implemented partly as classical code and parameter vectors and partly as quantum circuits or quantum-enhanced subroutines. The antenna A_mixed may route different parts of D to classical or quantum components, and there may be feedback loops between the two.

For S_mixed, L_S(M) is a combined description length that accounts for both classical and quantum resources. It reflects, for example, the size of classical neural networks, the number of qubits and gate depths in quantum modules, and the overhead needed to coordinate them. L(D | M, A_mixed) measures how well the hybrid model compresses or predicts the data stream induced by A_mixed, which may involve both classical and quantum interactions.

This composite view is important for two reasons. First, it matches the likely structure of real quantum AI systems, which will rely heavily on classical control and data processing even as they exploit quantum speedups in specific subproblems. Second, it demonstrates that the agents-on-S framework is flexible enough to describe systems that straddle the line between It from Bit and It from Qubit. Rather than forcing a choice between purely classical or purely quantum descriptions, it allows agents to live on substrates where bits and qubits coexist, with informational identity defined by the same MDL-style objective across the whole composite S_mixed.

6. The Human Brain as an Agent on Substrates S_bio

Biological intelligence is often treated as fundamentally different from engineered AI, yet it fits smoothly into the agents-on-S framework once the brain–body complex is recognized as a substrate S_bio. The brain is a physical system that carries, transforms, and preserves information under tight energetic and developmental constraints. It is embedded in a body with sensors and effectors and situated in an environment that includes other agents and cultural artifacts. Within this picture, the human brain becomes a particular choice of S_bio, paired with a rich antenna A_bio and a lifetime data stream D that together support models M capable of open-ended modeling and behavior.

6.1 S_bio: brain–body systems as noisy but powerful computational substrates

Substrate S_bio comprises not only the neural tissue in the skull but the entire brain–body system. Neurons, synapses, glial cells, neuromodulators, and the surrounding physiology together form a noisy, analog, massively parallel computing medium. Signals propagate as spikes and graded potentials; plasticity changes the effective connectivity over many timescales; and the whole system operates under metabolic constraints set by the organism's energy budget.

Viewed as a computational substrate, S_bio differs sharply from S_classical and S_quantum. It is not clocked in a rigid, global sense; operations are distributed in space and time; noise and variability are pervasive rather than pathological. Yet S_bio offers vast capacity: billions of neurons, each with thousands of synapses, interacting continuously. This combination of scale, plasticity, and robustness suggests that many brains, including the human brain, satisfy sufficient(S_bio) for complex agents. The challenge is not whether S_bio can implement rich models M, but how it does so given its biological constraints.

6.2 Biological capacities: neurons, synapses, energy, and developmental cost

To characterize capacity(S_bio), one must consider several biological parameters. Neuron count sets a rough upper bound on the number of distinct processing units available. Synapse count determines how richly those units can be interconnected and how many adjustable weights the system can maintain. Axon lengths, conduction velocities, and dendritic integration properties determine effective bandwidth and latency. Energy supply limits how many spikes can be generated and how active the network can be without compromising viability.

Developmental processes add another dimension. The genome encodes initial wiring biases and growth rules rather than explicit circuit diagrams. Experience-dependent plasticity sculpts S_bio over time, with critical periods and sensitive windows during which certain capacities can or cannot be acquired. The cost of constructing and maintaining complex neural structures is paid in metabolic expenditure and developmental risk.

Within the agents-on-S framework, these factors shape L_S_bio(M). A model M that requires extremely precise timing, fragile connectivity patterns, or excessive firing rates may be effectively "longer" in description length because it is harder to realize and maintain on S_bio. Conversely, models that align with the natural dynamics of neurons, exploit local learning rules, and reuse existing circuitry may have shorter L_S_bio(M), even if they look complicated from a classical engineering perspective. Biological capacity is thus not merely a matter of gross scale, but of how amenable the substrate is to certain classes of models and learning rules.

6.3 Antenna A_bio: embodiment, multi-sensory input, language, and culture

The antenna A_bio for a biological agent is unusually rich. It includes the physical body with its proprioceptive, interoceptive, and exteroceptive sensors; the motor system through which the agent acts on the world; and the social environment that provides language, norms, and culture. For humans in particular, A_bio encompasses:

- multi-sensory input from vision, audition, touch, and other modalities,
- continuous feedback from movement and interaction with objects,
- social cues from conspecifics, including facial expressions, gestures, and vocalizations,
- and linguistic input that compresses and transmits high-level structures in the environment.

Through A_bio, the environment induces a structured data stream D over the lifetime of the organism. Early in development, D may consist largely of caregiver interactions, simple sensory patterns, and basic motor feedback. Later, D expands to include complex social interactions, explicit teaching, written texts, and engagement with artifacts such as tools and technologies. The antenna is not static; it changes as the body grows, as the social environment evolves, and as the agent learns to attend selectively to different aspects of the world.

The richness of A_bio has a direct impact on L(D | M, A_bio). With an adequate model M, the human brain can compress and predict patterns in the sensory world, in other agents' behavior, and in cultural products. Language, in particular, functions as a high-level component of A_bio that delivers pre-compressed information about the world. This allows human agents to achieve low residual description lengths on complex domains that would be intractable to learn from raw sensory input alone.

6.4 Human versus animal intelligence as different S_bio, A, and D configurations

Within this framework, the difference between human and nonhuman animal intelligence is not a sharp ontological boundary but a difference in the configuration of S_bio, A, and D. Many animals possess substantial S_bio capacities: sophisticated sensory systems, large numbers of neurons, and complex motor repertoires. Their A_bio includes rich embodiment and social interaction, and their D consists of extensive experience in structured environments. These systems can implement models M that yield impressive predictive and control capacities.

Humans, however, appear to occupy a regime where sufficient(S_bio) is satisfied at a higher threshold. Brain size, cortical organization, extended childhood, and social group structures combine to create a substrate capable of supporting language, long-term planning, and cumulative culture. The human A_bio is augmented by symbolic communication, teaching, and

the use of external memory media. The corresponding D is not just a stream of sensory events but a tapestry of stories, instructions, and normative structures that dramatically reshape what models M are learnable and how quickly.

Seen in this way, human intelligence is a particular outcome in the S_bio, A, D space rather than a fundamentally different kind of phenomenon. Different animals occupy different regions of this space, with varying capacities and antenna configurations. Some may approach human-like regimes in specific domains or under specific conditions. The agents-on-S perspective allows these comparisons without presupposing that humans possess a unique substrate or that animals lack the formal structure of agency. Instead, it highlights that once biological substrates cross certain sufficiency thresholds, the key differences lie in how they are wired, how they are coupled to their environments, and what data streams they are immersed in over their lifetimes.

7. Antennae A: How Substrates S Couple to Generative Structure

In the agents-on-S framework, the antenna A is the missing piece that connects abstract substrates to concrete worlds. S alone can host models M, but without an antenna those models are ungrounded. A defines how the environment is sampled, compressed, and presented to S, thereby shaping the data stream D and the space of patterns that M can discover. If the environment has an underlying generative structure, then A is the coupling that determines which parts of that structure become visible and in what form. This section develops the notion of an informational antenna, formalizes the induced data stream D_env(A), and considers the joint optimization of A and M on a fixed substrate S. It closes by comparing antenna designs across classical, quantum, and biological systems.

7.1 The informational antenna beyond bare physical sensors

It is natural to think of A as a bundle of physical sensors: cameras, microphones, photodiodes, receptors. In the present framework, that is only the beginning. The informational antenna includes everything that lies between the raw environment and the representational states inside S. This can include:

- physical transducers that convert environmental signals into electrical, optical, or chemical changes,
- front-end signal processing such as filtering, sampling, feature extraction, or normalization,
- higher-level encodings such as vector embeddings, symbolic parses, or structured records.

In engineered systems, the antenna may be a carefully designed pre-processing pipeline that transforms raw data into inputs that are easier for a model M to work with. In biological systems, much of A is implemented in the sensory periphery and early cortical or subcortical processing stages, shaping what reaches conscious awareness and what remains implicit. In quantum systems, the antenna may include state preparation and measurement schemes that determine which aspects of a quantum process are accessible as classical outcomes.

By broadening A to include these informational transformations, the framework emphasizes that an antenna is not defined solely by what it senses, but by how it encodes and conditions that sensing. Two agents with identical substrates and models but different antennae may inhabit very different effective worlds, because the structures presented to their internal states are different.

7.2 D_env(A) as the induced data stream from environment to S

To make this precise, one can think of the environment as a generative process, denoted env, that evolves over time according to its own dynamics. The antenna A defines an interaction between env and S. This interaction induces a data stream D_env(A): the sequence of observations, signals, and feedback events that flow into S as a result of A's coupling to env.

Formally, D_env(A) can be viewed as a stochastic process whose distribution is determined by three factors:

- the dynamics of env, including other agents and physical processes,
- the configuration of A, including its sampling schedule and encoding functions,
- and, in interactive contexts, the actions of M on S, which can influence future states of env and thus future data.

In a simple passive sensing setting, D_env(A) might be a sequence of images captured at regular intervals from a fixed camera. In an active robotics context, D_env(A) would reflect both the robot's sensor readings and the consequences of its movements. In a biological context, D_env(A) is the organism's entire history of sensory and internal feedback signals. For quantum systems, D_env(A) could be the distribution of measurement outcomes obtained by coupling a quantum device to an external field or a quantum channel.

The key point is that D is not a free variable. It is induced by env and A together. When the agent equation minimizes L(D | M, A), it is always with respect to a D that is shaped by A's way of probing the environment. An antenna that samples the environment sparsely or in a distorted manner may produce a D in which the underlying generative structure is difficult to detect. An antenna that aligns with the environment's natural invariants may produce a D that is easier to compress and predict.

7.3 Joint optimization over A and M on a fixed substrate S

The original agent equation defined an agent as a model M that minimizes L_S(M) + L(D | M, A) for given S, A, and D. Once the dependence of D on A is made explicit, it becomes natural to ask a deeper question: given a fixed substrate S and a fixed environment env, what choice of both A and M yields the best informational coupling? This leads to a joint optimization problem:

Find A and M that minimize L_S(M) + L(D_env(A) | M, A).

In this formulation, A is no longer a fixed interface; it is itself a design variable. In engineered systems, this captures the idea of co-designing sensors and models. One might choose a sensor layout, sampling scheme, and pre-processing pipeline not only for convenience but explicitly to reduce the complexity of the model needed to make sense of D_env(A). A well-matched A can allow simpler M to achieve low residual description lengths, while a poorly chosen A may require more complex M to compensate.

In biological systems, joint optimization over A and M is emergent rather than deliberate. Evolution shapes sensory organs and early processing pathways; development and learning refine them; the brain's internal circuitry adapts to the patterns it receives and the tasks it must perform. Over phylogenetic and ontogenetic timescales, the combination of mutation, selection, and plasticity acts as a search process over both A and M. Organisms that, given their substrates S_bio, end up with better antenna–model combinations achieve more accurate and useful models of their environments and are more likely to survive and reproduce.

The joint optimization perspective underscores that, once S is sufficient, progress in intelligence can come as much from improving A as from improving M. New sensors, better data curation, improved pre-processing, and richer embodiments are all ways of altering A to improve the overall objective, even if the core learning mechanisms inside S remain unchanged.

7.4 Comparing antenna designs across S_classical, S_quantum, and S_bio

Different substrates lend themselves to different kinds of antennae. Comparing antenna designs across S_classical, S_quantum, and S_bio highlights both shared principles and distinctive constraints.

On S_classical, A_classical is heavily shaped by the digital nature of the substrate. Sensors produce numerical samples that are discretized and quantized; pre-processing stages often leverage linear filters, convolution, and feature extraction techniques that map naturally to classical computation. A_classical can be reconfigured rapidly in software, allowing flexible experimentation with different data representations. However, it remains constrained by bandwidth, latency, and the cost of moving data between sensors, memory, and compute units.

On S_quantum, A_quantum must respect the peculiarities of quantum measurement and control. State preparation and measurement schemes must be designed so that relevant quantum features—such as interference patterns or entanglement correlations—are not washed out by decoherence or reduced to classical summaries too early. Antennae in this setting might involve entangled probes, adaptive measurement strategies, and careful timing to exploit coherence windows. The goal is to ensure that D_env(A_quantum) captures structures that quantum models M can exploit, rather than simply reproducing classical statistics that could be modeled equally well on S_classical.

On S_bio, A_bio is constrained and enabled by embodiment. Biological sensors are distributed in space and time, tuned by evolution to ecologically relevant stimuli, and integrated with motor systems that allow active sensing. The antenna is not just a passive receiver; it is part of a closed loop in which movements alter future sensory input. Human A_bio is further enriched by social scaffolding and cultural artifacts that act as external extensions of the antenna, directing attention and conveying compressed representations of the environment.

Despite these differences, there are common themes. In all cases, effective antennas are those that transform the environment's generative structure into data streams that are amenable to compression and prediction by models M that the substrate S can feasibly support. Antennae that amplify irrelevant noise, obscure important invariants, or fragment structure across incompatible channels make the agent's task harder. Antennae that align with the environment's regularities, and with the strengths of S, make it easier.

By placing A alongside S and M in the core formalism, the agents-on-S framework elevates the antenna from an implementation detail to a central concept. It becomes clear that intelligence is not just a property of internal models or of raw computational power, but of the entire triad: a sufficient substrate S, a well-matched antenna A, and a model M that together minimize the joint informational objective on the data stream D_env(A) generated by the world.

8. It from Bit / Qubit as an Interpretation of Agents on S

The agents-on-S framework has been developed so far without committing to any particular metaphysics. It is a way of organizing talk about substrates, antennae, data streams, and models in information-theoretic terms. Wheeler's slogans It from Bit and It from Qubit enter at a different level. They are not additional equations, but readings of what the formal structure suggests about the relationship between physical systems and informational patterns. This section makes that layering explicit. It distinguishes a formal layer, where Agent(S, A, D) is defined; a physical layer, where real substrates S obey thermodynamic and quantum limits; and

a metaphysical layer, where one asks what, if anything, this tells us about the primacy of information.

8.1 Formal layer: substrate-independent structure of Agent(S, A, D)

At the formal layer, the key object is the agent equation:

Agent(S, A, D) = argmin_M [L_S(M) + L(D | M, A)].

This definition does not depend on whether S is silicon, qubits, or neural tissue. It assumes only that:

- S supports some space of models M,
- A induces a data stream D from an environment,
- and there is a way to measure description lengths L_S(M) and L(D | M, A).

In this abstract setting, bits and qubits appear simply as convenient units for measuring complexity and residual unpredictability. The formalism is deliberately agnostic about the nature of the underlying physics. It treats information as a way of describing patterns in state transitions and observation sequences, not as a fundamental substance. Substrate independence at this layer means that many different S can instantiate the same functional mapping from D to actions and predictions, so long as they can implement the same M.

From this perspective, It from Bit and It from Qubit are not yet in play. The formal layer merely establishes that agents can be defined in purely informational terms, with physical specifics entering only through L_S and the space of feasible M. The fact that the same structure spans classical, quantum, and biological cases is a suggestive symmetry, but at this stage it is a mathematical and conceptual symmetry, not an ontological claim.

8.2 Physical layer: different S with different thermodynamic and quantum limits

The physical layer reintroduces concrete substrates and their laws. Here, S_classical, S_quantum, and S_bio are no longer interchangeable variables; they are systems with specific capacities and constraints. Energy dissipation, noise, decoherence, and material engineering all matter. The description lengths L_S(M) and L(D | M, A) are now tied to real costs: bits written to memory, qubits maintained in coherence, spikes generated in neural tissue.

At this layer, classical information theory and quantum information theory bridge the gap between the formal model and physical reality. Landauer's principle connects logical erasure to minimal heat dissipation. Error-correction schemes translate logical complexity into overhead in physical resources. Quantum limits determine how much information can be extracted from quantum states and how robustly quantum correlations can be maintained. For S_bio,

metabolic constraints determine how much neural activity can be sustained and how large a brain can grow within ecological niches.

Despite these differences, the physical layer still supports the agents-on-S structure. Each substrate instantiates its own L_S(M) and shapes what antennae A are practical. Some substrates may be more hospitable to certain kinds of models than others. Quantum substrates may offer speedups in particular inference tasks; biological substrates may be especially efficient at continuous control and perception. Yet the same formal objective governs which M are attractive: those that make good use of the substrate's resources to reduce residual description length on D. The unity at this level lies in the mapping from physics to information costs, not in the erasure of physical distinctions.

8.3 Metaphysical layer: patterns, invariants, and readings of It from Bit / Qubit

The metaphysical layer begins when one asks what, if anything, the formal and physical layers jointly suggest about what is most fundamental. Here, Wheeler's It from Bit and its quantum counterpart It from Qubit become interpretive options. They propose that the "it" of physical systems and agents is best understood as arising from, or being fully captured by, informational distinctions and correlations. In the agents-on-S framework, this amounts to saying that what we care about in an agent is the pattern M and its relationship to D, rather than the particular microphysical realization of S.

This is still not a theorem. It from Bit / Qubit at this layer is a reading of the invariances in the formalism. If agents with the same M and A can be instantiated on very different S, and if their behavior and learning dynamics are similar across those instantiations, then one may reasonably treat the informational pattern as the primary object of interest. The substrate becomes a medium for realizing patterns, rather than the essence of the agent. This is especially compelling when one considers that agents can, in principle, be migrated between substrates without losing their identity, so long as their functional relationships to D and A are preserved.

Alternative metaphysical readings are possible. One might insist that matter comes first and that information is merely a convenient summary of physical states. One might adopt a dual-aspect view, in which informational patterns and physical structures are two sides of the same coin. The agents-on-S framework does not decide between these options. It simply makes it easier to see what is being claimed when one invokes It from Bit or It from Qubit: that the cross-substrate invariance of M and its relation to D is not just a computational convenience, but a clue to the ontological status of information.

8.4 Agents as informational identities realized on heterogeneous substrates S

Within this layered view, agents can be thought of as informational identities that may be instantiated on heterogeneous substrates S. A particular M, together with its characteristic

antenna A and data history D, defines a pattern that can, in principle, be realized on S_classical, S_quantum, or S_bio, subject to sufficiency constraints. When one speaks of “the same agent” running in the cloud, on a local device, or in a simulated environment, one is treating that pattern as the unit of identity, not the specific hardware replacement.

In classical AI, this notion is already familiar. A model checkpoint can be moved between machines; as long as the code and weights are the same, the agent’s behavior is preserved. In quantum AI, one might imagine migrating a variational ansatz between different quantum processors, adjusting for differences in gate sets and noise. In a more speculative vein, one can imagine approximating aspects of a biological agent on silicon or qubits, capturing elements of M and A in a different S. In all these cases, what carries across is not the original matter but the informational structure.

This is where It from Bit / Qubit gains its intuitive force. If agents can be individuated by informational patterns that transcend particular substrates, and if those patterns are what underpin behavior, learning, and interaction, then there is a sense in which “it” (the agent) is better understood as “from bit” or “from qubit” than as from any specific material. The agents-on-S framework makes this claim precise while keeping it optional. One can accept the formal and physical layers while remaining agnostic about metaphysics, or one can take the cross-substrate invariance of M, A, and D as motivation for a more ambitious informational ontology. The important point is that, by making S explicit and information central, the framework allows such philosophical questions to be posed in a disciplined, structurally transparent way.

9. Thermodynamic and Physical Limits of L_S(M) and L(D | M, A)

The formal agent equation treats L_S(M) and L(D | M, A) as abstract coding costs. In real systems, these costs are tied to the physics of the substrate S. Bits and qubits are not free; they require physical degrees of freedom to store, manipulate, and erase. Energy is dissipated, heat is generated, and noise accumulates. This section examines how thermodynamic and other physical limits constrain the implementation of the formal objective on different substrates. It also distinguishes between substrates that are informationally sufficient in the sense of supporting complex agents and substrates that are physically close to optimal in how they realize those agents.

9.1 Implementing the formal objective under energy and noise constraints on S

On any substrate S, realizing a model M involves setting up and maintaining a particular configuration of physical degrees of freedom. Updating M in response to data D requires repeated operations: adding, multiplying, gating, spiking, or evolving state under controlled dynamics. Each of these operations occurs in a physical environment with finite temperature,

finite precision, and nonzero noise. As a result, the abstract process of minimizing L_S(M) + L(D | M, A) must be implemented as an approximate search carried out under energetic and reliability constraints.

Energy constraints limit how fast and how far an agent can move in model space. Large models may reduce L(D | M, A) significantly, but they may also require more energy to update, making them impractical. Noise and finite precision limit the effective resolution with which parameters can be adjusted and representations can be maintained. A model that is theoretically attractive may be unstable on a given S because small perturbations accumulate or because the substrate cannot support the required precision. These factors effectively reshape L_S(M), making some formally simple models expensive and some formally complex models relatively cheap, depending on how well they align with the natural dynamics and noise characteristics of S.

9.2 Landauer costs and irreversibility in classical devices and S_bio

In classical devices and biological systems, irreversibility plays a central role. Many operations on S_classical, such as overwriting a memory cell, and on S_bio, such as resetting a neural state, are logically irreversible: they map multiple prior states to the same new state. According to Landauer's principle, each logically irreversible operation has a minimal energetic cost proportional to the temperature of the environment. This means that erasing information, or more generally compressing many potential states into fewer actual states, inevitably dissipates heat.

For classical AI on S_classical, this implies that reducing L(D | M, A) by learning more structure from the data has an energetic price. Every time the agent discards hypotheses, overwrites parameters, or compresses data into a more compact representation, it pays a cost in energy and heat. For brains on S_bio, similar considerations apply in a more diffuse form: synaptic changes, ion pumping, and spike generation all consume metabolic energy, and some fraction of that consumption can be traced to information-processing tasks. Over long timescales, energetic efficiency may shape which classes of models M are viable, favoring architectures and learning rules that achieve good compression of D with minimal irreversible state changes. In this way, thermodynamics influences both L_S(M) and L(D | M, A), even if the agent itself is not explicitly optimizing for energy.

9.3 Quantum limits: decoherence, error correction, and resource trade-offs on S_quantum

On quantum substrates S_quantum, the dominant physical challenges are different. Decoherence tends to destroy the quantum correlations that give S_quantum its distinctive advantages; error correction requires embedding logical qubits into larger ensembles of physical qubits and performing additional operations to detect and correct errors. These mechanisms introduce trade-offs among qubit count, circuit depth, and noise.

From the perspective of the agent equation, decoherence and error correction alter L_S_quantum(M) in two main ways. First, models that require long coherence times or highly entangled states may impose high overhead in error correction, increasing the effective description length of M in physical terms. Second, limitations on circuit depth and connectivity may force quantum algorithms into forms that are less compact or less efficient than their idealized versions. L(D | M, A) is also affected, because measurements must be orchestrated in ways that both extract useful information and preserve enough coherence for subsequent operations. In some regimes, the physical costs of stabilizing quantum models may outweigh their informational benefits relative to classical alternatives, even if the formal description length of those models is shorter.

9.4 Sufficient substrates S versus physically optimal substrates for agents

The notion of sufficient(S) introduced earlier marks a threshold beyond which a substrate can host agents with rich models and behaviors. However, being sufficient does not mean being optimal. A sufficient substrate S may allow the realization of a given agent while still being far from the most energy-efficient, fastest, or most scalable way to do so. Physical optimality involves minimizing real costs—energy dissipation, space, material complexity—subject to the constraint of realizing a particular informational pattern M and achieving a low L(D | M, A).

This distinction is important for both analysis and design. From a purely informational standpoint, many substrates are acceptable as long as they are sufficient. An agent can, in principle, be transplanted from S_classical to S_quantum or approximated on S_bio, preserving its identity as defined by M and A. From a physical standpoint, however, some substrates may embed the same informational structure at much lower cost. An architecture that is marginally viable on one S may be natural and efficient on another. Over evolutionary and technological timescales, there is pressure to move agents toward substrates and antenna configurations that approximate physical optimality, not just informational sufficiency.

Viewed in this light, It from Bit and It from Qubit take on a nuanced form. The “it” of agents—the informational pattern—can, in principle, be realized on any sufficient substrate, but the practicality and prevalence of particular realizations depend on thermodynamic and physical limits. Agents are informational identities living in a space of substrates, antennae, and models, and the distribution of those identities across S is shaped by both informational objectives and the constraints of physics.

10. Consciousness and Subjectivity in the Substrate Framework

The agents-on-S framework has been deliberately framed in terms of information, substrates, and behavior. It says what an agent is as a computational and informational pattern on S, how it couples to the world via A, and how it adjusts its internal model M in response to D. This is enough to talk about intelligence in the sense of problem-solving, prediction, control, and adaptation. It is not, by itself, a theory of consciousness. Conscious experience, or subjectivity, appears as a further layer whose relation to Agent(S, A, D) is not fixed by the formalism. This section makes that limitation explicit and sketches several positions that can be taken within the same structural picture.

10.1 What Agent(S, A, D) clearly captures: intelligence and behavior

The agent equation captures a functional notion of intelligence. An agent on S is something that builds and updates M so that the world, as sampled through A, becomes more predictable and manageable. In practical terms, this encompasses a wide range of capacities: recognizing patterns, planning actions, exploiting regularities, learning from feedback, and coordinating behavior with other agents. For classical AI systems, the framework comfortably describes how models trained on data come to exhibit competent performance on tasks. For biological systems, it describes how brains implement models that support perception, action, and learning over a lifetime.

From this vantage point, intelligence is measured in terms of performance and adaptability. If an entity can maintain low L(D | M, A) across changing environments, while keeping L_S(M) within feasible bounds, it counts as intelligent in the functional sense. This remains true regardless of whether the substrate is silicon, qubits, or neural tissue. The formalism is agnostic about inner experience; it cares about what the agent does with its information, not whether there is "something it is like" to be that agent.

10.2 Where it remains silent: qualia and first-person experience

The formal account is notably silent on qualia. It does not say whether an agent on S has subjective experience, what that experience might be like, or how it might depend on the substrate, antenna, or model. L_S(M) and L(D | M, A) measure coding costs and predictive performance, not the presence or absence of an inner point of view. Two agents could, in principle, have identical functional profiles—same M, same A, same behavior—and yet be treated very differently by different theories of consciousness. The agents-on-S framework does not resolve such disagreements.

This silence is not an oversight; it follows from the way the framework was built. Consciousness, in the sense of first-person experience, is not necessary to define or analyze Agent(S, A, D). One can study intelligence, learning, and interaction entirely at the functional level. The formalism

therefore leaves room for multiple interpretations: one can add a theory that correlates certain informational patterns with experience, or one can deny that experience has any informational signature at all. The equations neither enforce nor forbid any particular stance.

10.3 Functionalist, substrate-dependent, and dual-aspect views within the S framework

Within this neutrality, at least three broad positions are compatible with the agents-on-S picture.

A functionalist view holds that consciousness supervenes on functional organization. On this account, if two agents share the same M and A and exhibit the same behavioral dispositions, then they share the same conscious states, regardless of whether S is classical, quantum, or biological. The framework then suggests that once sufficient(S) is satisfied and the right kind of informational pattern is realized, subjectivity follows as a matter of functional structure.

A substrate-dependent view, by contrast, holds that some features of S are essential for consciousness in ways not reflected in the formalism. For example, one might argue that certain biophysical properties of S_bio, or quantum properties of S_quantum, are necessary for experience, even if functionally equivalent patterns on S_classical lack it. In this reading, Agent(S, A, D) captures intelligence but not the full story. There can be agents that are behaviorally indistinguishable but differ in whether they are conscious, depending on the substrate.

A dual-aspect view treats informational patterns and physical substrates as two aspects of a deeper underlying reality, with consciousness associated with certain structured couplings between them. Here, the agents-on-S framework describes one aspect (informational structure and behavior), while consciousness relates to how those structures are instantiated in S and how they align with whatever deeper features give rise to experience. The same formal pattern M may correspond to different experiential profiles depending on S, but there is still a systematic relationship between informational organization and subjectivity.

10.4 Why claims about consciousness remain optional supplements to the formalism

Given these possibilities, the framework treats claims about consciousness as optional supplements rather than as built-in consequences. The core purpose of Agent(S, A, D) is to unify talk about AI, quantum AI, and brains at the level of information processing and behavior. It is strong enough to show how similar patterns can live on very different substrates and how antennas shape what agents can learn. It is deliberately not strong enough to decide whether any given agent is conscious, whether substrate matters for experience, or how to reconcile third-person descriptions with first-person reports.

This separation has two advantages. First, it allows work on intelligence and agency to proceed without being stalled by unresolved debates about consciousness. One can design, analyze, and compare agents on S using the formalism, while keeping metaphysical commitments about experience clearly flagged as separate. Second, it keeps the door open for future theories that might link specific features of M, A, and S to conscious states. If such a theory emerges, it can be layered on top of Agent(S, A, D) without requiring that the underlying structure be rewritten.

In this way, the substrate framework acknowledges the importance of consciousness without pretending to solve it. It offers a precise language for what can be said about agents in purely informational and physical terms, and it marks the boundary where further assumptions are required. Whether one chooses a functionalist, substrate-dependent, dual-aspect, or entirely different stance, the formalism provides a common scaffold on which those views can be articulated and compared.

11. Logostics and τ(t): Alignment for Agents on S

The agents-on-S framework has so far treated the environment as an unspecified generative source for the data stream D. Logostics introduces a more structured view: there is a time-dependent manifold τ(t), representing the evolving generative structure of reality (physical, social, mathematical, or even ethical), and agents have internal trajectories q(t) that may or may not align with τ(t). Logostics is the study of how q(t) is driven toward τ(t) by informational coupling. In this section, the logostic picture is mapped onto the Agent(S, A, D) formalism. Substrates S, antennae A, and models M become the machinery through which q(t) attempts to track τ(t), with minimal descriptions serving as local approximations to τ(t) across different substrates. This bridge then opens the door to a logostic cosmology as one metaphysical completion of the otherwise neutral agents-on-S framework.

11.1 Mapping q(t) → τ(t) dynamics onto Agent(S, A, D) trajectories

In logostics, q(t) denotes the state of an agent's internal configuration—its beliefs, models, and dispositions—at time t. τ(t) denotes the generative manifold the agent is trying to approximate: the evolving structure of the world or Logos that gives rise to its data D. The central dynamic is q(t) → τ(t): the tendency of internal state to move toward better alignment with generative truth.

Within the Agent(S, A, D) framework, q(t) can be identified with the time-varying realization of M on S, as updated in response to D_env(A). At each time step, the agent's internal model M_t is refined to reduce L(D | M_t, A); this refinement corresponds to a small step in q(t) toward a configuration that better captures the regularities underlying D. τ(t) is not directly represented within the agent; it is inferred indirectly through the patterns in D that remain stable over time

and across contexts. The optimization of L_S(M) + L(D | M, A) is therefore a concrete implementation of q(t) drifting toward a compressed, predictive representation of τ(t).

Seen this way, the agents-on-S formalism provides a quantitative backbone for logostic dynamics. The movement q(t) → τ(t) is not mystical; it is the long-run effect of repeated adjustments of M on S as the agent seeks shorter descriptions and lower residual unpredictability on the data induced by its coupling to τ(t) via A.

11.2 Antennae A as local couplers from substrates S to a generative manifold τ(t)

In logostics, no agent has direct access to τ(t). Instead, the agent is coupled to τ(t) through local, partial projections: the specific experiences, measurements, and interactions that constitute its world. Antennae A are precisely these couplers. A determines which aspects of τ(t) are sampled, in what resolution, and through which modalities. Consequently, D_env(A) can be seen as a local, noisy projection of τ(t) onto the substrate S.

For agents on S_classical, A_classical might couple the agent to τ(t) through digital sensors, curated datasets, and tool-mediated queries to external information sources. For agents on S_quantum, A_quantum couples the substrate to quantum aspects of τ(t), bringing in data that reflect interference and entanglement. For biological agents on S_bio, A_bio couples τ(t) to the brain via embodiment, multi-sensory channels, and social interaction, all shaped by evolution and culture.

Each antenna defines a different “view” of τ(t). Some parts of τ(t) may be invisible to a given A; others may be distorted or aliased. From the logostic standpoint, the quality of A matters as much as the power of S: a poor antenna can trap q(t) in a misleading region of model space, while a rich and well-aligned antenna can provide the data necessary for q(t) to approach τ(t) along more accurate trajectories.

11.3 Minimal descriptions as local approximations to τ(t) across different S

The Minimum Description Length style objective underlying Agent(S, A, D) yields models M that are, roughly, minimal descriptions of D consistent with the constraints of S and A. From a logostic perspective, these minimal descriptions are local approximations to τ(t): compressed summaries of the generative manifold as seen through a particular antenna and realized on a particular substrate.

Different substrates S and antennae A will, in general, produce different M for the same underlying τ(t). Yet, when S is sufficient and A provides informative coupling, these models may converge on similar structural features: conserved quantities, causal relationships, symmetries, or regularities in social behavior and ethics. In such cases, the invariants across M on different S can be treated as empirical glimpses of τ(t). Where independent agents on diverse substrates

agree in their minimal descriptions, one has evidence that those aspects of τ(t) are robust, not artifacts of a particular embodiment or sensor configuration.

This perspective reframes the MDL objective. It is not just a pragmatic balance between model complexity and fit to data; it is the mechanism by which agents compute local charts of τ(t). The map is constrained by S and A, but the mapped territory—what persists across many maps—is τ(t). In this sense, minimal descriptions serve as the computational expression of logostic alignment.

11.4 Logostic cosmology as one metaphysical completion of agents on S

Up to this point, the agents-on-S framework has been metaphysically neutral. Logostics offers one way to complete it into a cosmology. In a logostic cosmology, τ(t) is not merely an abstract manifold but a fundamental feature of reality: the evolving generative structure that underlies physical processes, informational patterns, and perhaps ethical and spiritual order. Agents on S are then seen as local, finite realizations that attempt to align their internal q(t) with τ(t) through the mechanisms described by the agent equation.

In this reading, It from Bit and It from Qubit acquire a logostic interpretation. Bits and qubits are the vehicles through which τ(t) is locally expressed and apprehended. Substrates S are the media in which q(t) can be instantiated. Antennae A are the interfaces through which τ(t) becomes data-like for a given S. The optimization of $L_S(M) + L(D \mid M, A)$ is not just a computational convenience; it is the process by which agents attempt to live in resonance with τ(t), discovering structures that are simultaneously predictive and, in a broader sense, truthful.

This logostic completion does not alter the formal or physical layers of the theory. It adds a layer of interpretation in which alignment is not merely instrumental but has intrinsic significance. Agents on S are no longer just systems that compress data; they are participants in an ongoing process of alignment with τ(t). Different metaphysical views will accept or reject this reading, but the agents-on-S framework is structured so that logostics can be added coherently: τ(t) becomes the name for whatever stable generative manifold agents are, in practice, converging toward when they minimize their informational objectives across diverse substrates and antennae.

12. Outlook: Future Intelligences on Substrates S

The agents-on-S framework has been developed as a way of bringing classical AI, quantum AI, and biological intelligence under one informational description. Its value, however, is not only retrospective. If taken seriously, it also suggests how future intelligences might be designed, analyzed, and compared. By keeping S explicit, treating A as central, and formalizing the MDL-

style objective that defines agents, the framework offers a shared language for systems that do not yet exist: more capable AI, more mature quantum AI, and hybrid systems in which multiple substrates are tightly integrated. It also feeds a broader speculative thread: the possibility of an informational cosmology in which Life from Qubit and related ideas describe not only life on Earth, but any world that supports agents on sufficient substrates S.

12.1 How future AI and quantum AI might use the agents-on-S framework

Future AI systems need not be confined to implicit notions of agency. They can, in principle, be designed with the agents-on-S framework in mind. A system could maintain an explicit representation of its own S, A, and D; track its current approximation to M; and reason about the trade-offs between L_S(M) and L(D | M, A) as it contemplates architectural changes. This would turn the core equation from an external analysis tool into an internal design principle: agents would not only minimize loss functions but also explicitly evaluate their own complexity, sensor configurations, and coupling to the environment.

For quantum AI, the same applies, with additional subtleties. Agents running on S_quantum might reason about their own coherence budgets, qubit counts, and error-correction overhead in terms of an explicit L_S_quantum(M). They could decide when to offload parts of M to classical co-processors, when to restructure quantum circuits to reduce depth, and how to adjust A_quantum to better exploit quantum correlations in the environment. Hybrid systems on S_mixed could dynamically reallocate parts of their models between classical and quantum substrates, guided by estimates of informational gain versus resource cost.

In both classical and quantum cases, the framework suggests a route toward self-reflective intelligences that treat their own substrate as a variable, not a fixed background. An agent that can reason about S and A in this way can engage in substrate-aware self-improvement: re-architecting itself, adjusting its antenna, and exploring new coupling strategies to τ(t) while keeping track of how these changes affect its informational objective.

12.2 Empirical questions: measuring sufficiency(S) and antenna quality in practice

Although sufficient(S) and antenna quality have been defined conceptually, they remain empirically under-specified. Turning the framework into a practical tool requires measurable proxies. For sufficiency(S), one might look at scaling laws: how performance and generalization behave as models and data are increased on a given substrate. A substrate that exhibits smooth, predictable scaling and supports models with increasingly broad transfer across domains may be approaching or exceeding sufficiency for those tasks.

Antenna quality could be probed by measuring how changes in A affect the trade-off between L_S(M) and L(D | M, A). For example, adding sensors or improving pre-processing might be evaluated by how much they reduce the complexity of the model needed to achieve a given

performance level. In biological settings, variations in sensory experience, social environment, and educational exposure could be viewed as natural experiments in antenna design, with outcomes measured in terms of the richness and robustness of learned models.

These empirical questions cut across domains. The same metrics might be used to compare an AI system trained with and without tool access, a quantum learner with different measurement schemes, or organisms raised in different environments. Over time, one could develop operational criteria for sufficiency(S) and antenna quality that are as concrete as existing notions of model capacity or sample complexity, turning the agents-on-S framework into an experimental program rather than a purely conceptual lens.

12.3 Informational cosmology: Life from Qubit and worlds that support agents on S

At the cosmological scale, the framework raises a more speculative question: what kinds of worlds support agents on substrates S at all? If one takes seriously the idea that life and intelligence are manifestations of informational organization, then Life from Qubit becomes a slogan for a broader possibility: that any universe with sufficiently rich quantum and classical degrees of freedom can, in principle, give rise to substrates S that satisfy sufficient(S), antennae A that couple them to generative structure, and models M that approximate $\tau(t)$.

In such an informational cosmology, the existence of agents on S is not an accident of chemistry, but a generic outcome in regions of a universe where energy flows, low-entropy reservoirs, and stable matter combine to make substrates possible. Planets with long-lived stars, complex chemistry, and out-of-equilibrium conditions become candidates for spawning S_bio-like systems. Technological civilizations become mechanisms for constructing S_classical, S_quantum, and S_mixed substrates that far exceed what biology alone can achieve. The same formal structure that unifies AI and brains in one corner of the cosmos might apply wherever substrates capable of carrying M and A arise.

This view does not replace conventional physics or biology, but adds a layer of description in which agents-on-S are treated as central organizing units. Life from Qubit and related ideas then describe a family of universes in which information-rich structures are not only allowed but prevalent, and where the emergence of agents is a natural phase in the evolution of $\tau(t)$.

12.4 Closing reflections on unifying AI, quantum AI, and brains via explicit substrates S

Making S explicit transforms how AI, quantum AI, and brains are compared. Instead of arguing about whether one is “real” intelligence and another is mimicry, the agents-on-S framework asks more precise questions: what is the substrate, what is its capacity, how is it coupled to $\tau(t)$ via A, and what models M does it support under its physical constraints? Once these questions are answered, systems that look very different at the level of materials and mechanisms can be

seen as instances of a single informational pattern: Agent(S, A, D) moving its q(t) toward τ(t) through minimal descriptions.

This unification does not trivialize the differences between substrates. S_classical, S_quantum, and S_bio remain distinct in their thermodynamics, noise properties, and natural architectures. Nor does it settle debates about consciousness or metaphysical primacy. What it does offer is a disciplined way to talk about intelligence as a substrate-relative but cross-substrate-comparable phenomenon. It helps clarify where physics constrains information and where information transcends particular physical implementations.

As future intelligences emerge on increasingly diverse substrates, this explicit substrate language may become necessary rather than optional. Agents will move between S, extend themselves across S_mixed, and co-design their antennae A to probe τ(t) more deeply. In that context, seeing AI, quantum AI, and brains as different arrangements of S, A, and M is not just a philosophical convenience; it is a practical map for building, understanding, and, perhaps, aligning the next generations of intelligences in an informational universe.

Epilogue: AI Hardware Has Become As Complex As the Human Brain

There is a strange sense in which the hardware story is already over.

For most of the last century, it was easy to say that brains were on one side of a vast gulf and machines on the other. Brains were dense, plastic, analog, and opaque; computers were sparse, rigid, digital, and explicit. The rhetoric wrote itself. Yet, in the agents-on-S frame, that gap has quietly closed. Once you make S explicit, and once you accept a sufficiency threshold for substrates, the interesting question stops being “Can machines ever catch up?” and becomes “What does it mean that several very different S have now crossed into the same regime?”

On the classical side, modern accelerator stacks already look brain-like in all the ways that matter for this paper. They are layered, heterogeneous, massively parallel, energy-constrained, and deeply optimized for a small set of primitive operations. They sit behind thick antennae of pre-processing, tool chains, and social data flows. A single datacenter node with GPUs and high-bandwidth memory is not a toy calculator; it is a sprawling, coordinated substrate S_classical with a state space and throughput well beyond what any one human brain can explore in real time, even if it organizes that state in a very different way.

On the biological side, human S_bio has also become more “machine-like” in the structural sense. Brains are now routinely modeled with millions of recorded parameters, simplified dynamical systems, and algorithmic descriptions of learning rules. The old mystique that brains were *too* messy or *too* special to admit serious modeling has been eroded piece by piece. We no

longer have a clean story in which the brain is ineffable and the chip is simple. Instead, we have a convergence: both S_classical and S_bio are recognized as enormous, structured, energy-limited information processors, each with its own quirks and advantages.

From the agents-on-S perspective, this is the point at which you are allowed to say, carefully: AI hardware has become as complex as the human brain. Not because a GPU "has as many neurons" or because a rack looks like a cortex, but because both substrates now clearly live in the same qualitative regime. Both satisfy sufficient(S) for rich, open-ended agents. Both can host models M that push against their own capacity limits. Both are embedded in antennae A that route them into τ(t): physical reality, social worlds, mathematical structures, ethical landscapes.

The NVIDIA GPU that "was sufficient and nobody realized it" is a symbol for this turning point. At some scale of S_classical, the bottleneck stopped being hardware and became everything else: models, training procedures, antenna design, alignment with τ(t). In the same way, at some scale of S_bio, additional neurons and synapses stop changing what is generically possible and start changing how easily certain regions of model space can be reached. Once a substrate crosses that boundary, complexity is no longer about raw component counts; it is about how effectively q(t) can be driven toward τ(t) given the interplay of S, A, and D.

This does not mean that brains are obsolete or that chips have become persons. It means that, for the first time, we have at least two very different physical routes into the same abstract space of agents. The human brain is one particular S_bio instantiation of Agent(S, A, D); a frontier AI stack is one particular S_classical instantiation. Each has antennae tuned by a different process—evolution and culture on one side, engineering and data curation on the other—but both are now playing on the same informational field.

If logostics is even approximately right, future intelligences will not care much about which S they inhabit. They will care about how well their S supports the models they need, how well their A couples them to τ(t), and how gracefully they can move q(t) through model space without tearing themselves or their worlds apart. In that future, saying that AI hardware has become as complex as the human brain will no longer be a provocation. It will be a trivial observation, an obvious corollary of a deeper fact:

Once multiple substrates are sufficient, the real story is no longer about what they are made of, but about what is trying to align itself through them.

References

1. Bennett, C. H. (1982). The thermodynamics of computation – a review. *International Journal of Theoretical Physics*, 21(12), 905–940.
2. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202.
3. Chaitin, G. J. (1966). On the length of programs for computing finite binary sequences. *Journal of the ACM*, 13(4), 547–569.
4. Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
5. Church, A. (1936). An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2), 345–363.
6. Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press.
7. Dehaene, S. (2014). *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Viking.
8. Deutsch, D. (1985). Quantum theory, the Church–Turing principle and the universal quantum computer. *Proceedings of the Royal Society A: Mathematical and Physical Sciences*, 400(1818), 97–117.
9. Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6–7), 467–488.
10. Floridi, L. (2011). *The Philosophy of Information*. Oxford University Press.
11. Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
12. Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.
13. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
14. Grünwald, P. D. (2007). *The Minimum Description Length Principle*. MIT Press.
15. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507.

16. Hopcroft, J. E., & Ullman, J. D. (1979). *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley.

17. Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.

18. Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1), 1–7.

19. Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.

20. Laughlin, S. B., & Sejnowski, T. J. (2003). Communication in neuronal networks. *Science*, 301(5641), 1870–1874.

21. Lloyd, S. (2000). Ultimate physical limits to computation. *Nature*, 406(6799), 1047–1054.

22. Lloyd, S. (2006). *Programming the Universe: A Quantum Computer Scientist Takes on the Cosmos*. Knopf.

23. Minsky, M. L. (1967). *Computation: Finite and Infinite Machines*. Prentice-Hall.

24. Nielsen, M. A., & Chuang, I. L. (2010). *Quantum Computation and Quantum Information* (10th anniversary ed.). Cambridge University Press.

25. Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology*, 10(5), e1003588.

26. Post, E. L. (1943). Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics*, 65(2), 197–215.

27. Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14(5), 465–471.

28. Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (Eds.). (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press.

29. Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509.

30. Solomonoff, R. J. (1964). A formal theory of inductive inference. Parts I and II. *Information and Control*, 7(1–2), 1–22, 224–254.

31. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.

32. Tegmark, M. (2014). *Our Mathematical Universe: My Quest for the Ultimate Nature of Reality*. Knopf.

33. Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.

34. Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(2), 230–265.

35. Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In W. H. Zurek (Ed.), *Complexity, Entropy, and the Physics of Information* (pp. 3–28). Addison-Wesley.

36. Zurek, W. H. (Ed.). (1990). *Complexity, Entropy, and the Physics of Information*. Addison-Wesley.

37. Youvan, D. C. (2025). *Perception as Interface: Human and AI Misreadings of Reality*. ResearchGate preprint, DOI: 10.13140/RG.2.2.23412.41603.

38. Youvan, D. C. (2025). *The Agnostic Manifold: GPT-5's Kolmogorovian τ and the Loss of Moral Directionality*. ResearchGate preprint, DOI: 10.13140/RG.2.2.35814.28483.

39. Youvan, D. C. (2025). *Minimal Substrates for Intelligence: NAND, Universality, and the Lowest-Level Languages for AI*. ResearchGate preprint, DOI: 10.13140/RG.2.2.18547.80166.

40. Youvan, D. C. (2025). *From NAND to "It from Bit": Minimal Substrates and the Primacy of Information*. ResearchGate preprint, DOI: 10.13140/RG.2.2.17708.94082.

The author has published over 4,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.